

CLAIMS

What is claimed is:

- 5 1. In a multiprocessing computer system comprising a plurality of
processing nodes interconnected through an interconnect structure, wherein the
plurality of processing nodes includes a first processing node, a second processing
node, and a third processing node, a method for lock request arbitration comprising:
the first processing node transmitting a lock request to the second processing
10 node;
the second processing node determining whether the lock request is ready for
service; and
the second processing node issuing a broadcast message to the first and the
third processing nodes in response to determining that the lock request
15 is ready for service.
2. The method as recited in claim 1, wherein determining whether the lock
request is ready for service comprises:
the second processing node placing the lock request in a queue; and
20 the second processing node monitoring the queue to determine whether the
lock request is ready for service.
3. The method as recited in claim 2, wherein monitoring the queue
comprises:
25 sequentially processing each preceding lock request in the queue.
4. The method as recited in claim 1, wherein the broadcast message informs
the third processing node of servicing of the lock request.
- 30 5. The method as recited in claim 4, wherein the broadcast message includes
a lock bit to inform the third processing node of the servicing of the lock request.
6. The method as recited in claim 1, further comprising:

the first processing node sending a first probe response to the second processing node in response to the broadcast message; and the third processing node sending a second probe response to the second processing node in response to the broadcast message.

5

7. The method as recited in claim 6, wherein the third processing node sends the second probe response when the third processing node ceases issuance of new requests.

10

8. The method as recited in claim 6, further comprising: the second processing node informing the first processing node of availability of lock ownership upon receiving the first and the second probe response messages.

15

9. The method as recited in claim 8, wherein the second processing node informing the first processing node comprises: the second processing node transmitting a target done message to the first processing node to indicate the lock ownership.

20

10. In a multiprocessing computer system comprising a plurality of processing nodes interconnected through an interconnect structure, wherein the plurality of processing nodes includes a first processing node having ownership of a lock, a second processing node, and a third processing node, a method whereby the first processing node releases the ownership of the lock comprising:

25

the first processing node transmitting a lock release request to the second processing node; and the second processing node issuing a broadcast message to the first and the third processing nodes in response to the lock release request.

30

11. The method as recited in claim 10, wherein the broadcast message informs the third processing node of completion of lock operations within the multiprocessing computer system.

004080 080400 09633087

12. The method as recited in claim 11, wherein the broadcast message includes a lock bit to inform the third processing node of the completion of the lock operations.

5 13. The method as recited in claim 10, further comprising:
each of the first and the third processing nodes sending a first corresponding probe response message to the second processing node in response to the broadcast message.

10 14. The method as recited in claim 13, further comprising:
the second processing node transmitting a corresponding target done message to the first processing node upon receiving the first corresponding target done messages from the first and the third processing nodes.

15 15. A multiprocessing computer system comprising:
a plurality of processing nodes interconnected through an interconnect structure to each other and to shared system resources, wherein the plurality of processing nodes comprises:
a first processing node configured to generate a lock request prior to
20 commencing execution of an operation required access to the shared system resources;
a second processing node; and
a third processing node configured to receive and process the lock request and transmit a first broadcast message to the second
25 processing node in response to determining the lock request to be ready for service.

16. The multiprocessing computer system as recited in claim 15, wherein the third processing node comprises a queue for pending lock requests, wherein the third
30 processing node is configured to place the lock request in the queue and to monitor the queue to determine when the lock request is ready for service.

17. The multiprocessing computer system as recited in claim 15, wherein the first broadcast message comprises:

a lock bit, wherein the third processing node is configured to set the lock bit in response to determining the lock request ready for service.

18. The multiprocessing computer system as recited in claim 15, wherein the
5 interconnect structure includes a plurality of dual-unidirectional links.

19. The multiprocessing computer system as recited in claim 18, wherein
each dual-unidirectional link in the plurality of dual-unidirectional links interconnects
a respective pair of processing nodes from the plurality of processing nodes.
10

20. The multiprocessing computer system as recited in claim 18, wherein
each dual-unidirectional link in the plurality of dual-unidirectional links performs
packetized information transfer.

21. The multiprocessing computer system as recited in claim 15, wherein
15 each of the plurality of processing nodes includes:

a plurality of circuit elements comprising:

a processor core,

a cache memory, and

20 a memory controller; and

a plurality of interface ports, wherein each of the plurality of circuit
elements is coupled to at least one of the plurality of interface ports.

22. The multiprocessing computer system as recited in claim 21, wherein the
25 plurality of circuit elements further includes:

a bus bridge;

a graphics logic;

a bus controller; and

a peripheral device controller.

30

23. The multiprocessing computer system as recited in claim 21, wherein at
least one of the plurality of interface ports in the each of the plurality of processing
nodes is connected to a corresponding one of the plurality of dual-unidirectional links.

24. The multiprocessing computer system as recited in claim 15, further comprising:

a plurality of system memories; and

a plurality of memory buses, wherein each of the plurality of system

5 memories is coupled to a corresponding one of the plurality of processing nodes through a respective one of the plurality of system buses.

25. The multiprocessing computer system as recited in claim 15, wherein the
10 third processing node is configured to further transmit the first broadcast message to the first processing node.

26. The multiprocessing computer system as recited in claim 25, wherein
15 each of the first and the second processing nodes is configured to transmit a corresponding first probe response message to the third processing node in response to the first broadcast message.

27. The multiprocessing computer system as recited in claim 26, wherein the
20 second processing node is further configured to cease issuance of new requests and to transmit the corresponding first probe response message upon the cessation of the issuance.

28. The multiprocessing computer system as recited in claim 26, wherein the
25 third processing node is configured to transmit a first target done message to the first processing node after receiving the corresponding probe response messages from the first and the second processing nodes.

29. The multiprocessing computer system as recited in claim 28, wherein the
30 first processing node is configured to commence the execution of the lock operation after receiving the first target done message, and wherein the first processing node is further configured to transmit a lock release request to the third processing node after completion of the execution of the lock operation.

09633087.000400

30. The multiprocessing computer system as recited in claim 29, wherein the third processing node is configured to transmit a second broadcast message to the first and the second processing nodes in response to receiving the lock release request.

5 31. The multiprocessing computer system as recited in claim 30, wherein each of the first and the second processing nodes is configured to transmit a corresponding second probe response message to the third processing node in response to the second broadcast message.

10 32. The multiprocessing computer system as recited in claim 31, wherein the third processing node is configured to transmit a second target done message to the first processing node in response to receiving the corresponding second probe response messages from the first and the second processing nodes.

15 33. The multiprocessing computer system as recited in claim 32, wherein the third processing node is configured to place a pending lock request into service upon transmission of the second target done message, and the pending lock request is generated by one of the plurality of processing nodes and is stored within the third processing node.

20

004080" 0804100
09633087